2021 43rd Annual International Conference of the
IEEE Engineering in Medicine & Biology Society (EMBC)
Oct 31 - Nov 4, 2021. Virtual Conference

COVID-19 Volumetric Pulmonary Lesion Estimation on CT Images using a U-NET and Probabilistic Active Contour Segmentation

Leopoldo Cendejas-Zaragoza[1,2], Diomar E. Rodriguez-Obregon[3], Aldo R. Mejia-Rodriguez[3], Edgar R. Arce-Santana[3] and Alejandro Santos-Diaz[1,4]

***Abstract*— A two-step method for obtaining a volumetric estimation of COVID-19 related lesion from CT images is proposed. The first step consists in applying a U-NET convolutional neural network to provide a segmentation of the lung-parenchyma. This architecture is trained and validated using the Thoracic Volume and Pleural Effusion Segmentations in Diseased Lungs for Benchmarking Chest CT Processing Pipelines (PleThora) dataset, which is publicly available. The second step consists in obtaining the volumetric lesion estimation using an automatic algorithm based on a probabilistic active contour (PACO) region delimitation approach. Our pipeline successfully segmented COVID-19 related lesions in CT images, with exception of some mislabeled regions including lung airways and vasculature. Our workflow was applied to images in a cohort of 50 patients.**

I. INTRODUCTION

As of April 2021, COVID-19 pandemic has caused more than 3 million deaths worldwide [1]. The utility of Computed Tomography (CT) imaging has been recognized as an essential tool to diagnose the novel disease. Automatic algorithms have been proposed to differentiate COVID-19 disease from other lung-related illnesses. However, we hypothesize that relevant information can be extracted from this imaging technique to help elucidate patient prognosis. For this reason, we propose a workflow to automatically compute the volume related to a lesion due to COVID-19 in lung parenchyma.

II. METHODS

A. Study population and CT data acquisition

Data were collected retrospectively from 50 patients (18 female and 32 male, 55 $\pm$ 14 years) admitted to the intensive care unit at the INCMNSZ due to severe pneumonia caused by COVID-19 infection in the period from march to april of 2020. SARS-CoV-2 infection was confirmed by positive real time polymerase chain reaction (RT-PCR) test.This protocol was approved by the institutional ethics committee. Chest CT images were acquired using a GE Revolution EVO Gen 3 (GE Healthcare, Milwaukee, WI) system. The used imaging series consisted in a low resolution acquisition (70 - 101 slices) helical acquisition using a lung window, image size of 512x512 pixels, slice thickness of 1.25-3.75 mm and 1.25 mm spacing between slices. X-Ray tube voltage and current were set to 140 kV and 100-300 mA, respectively.

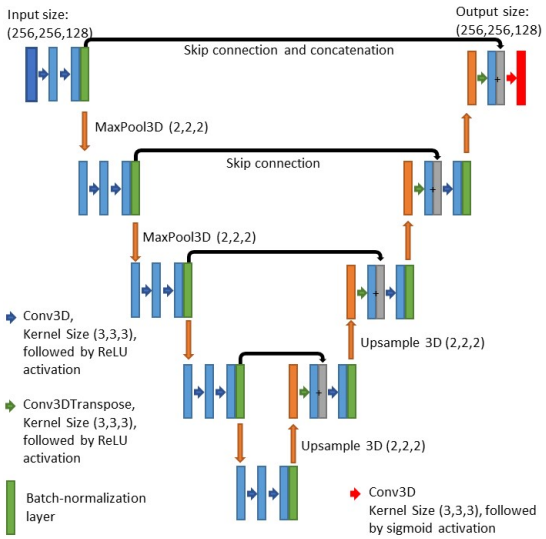

Fig. 1. U-NET architecture

B. Lung parenchyma segmentation

A modified U-NET architecture from that presented in [2] was proposed to extract the lung parenchyma from the raw CT images (see Fig. 1). The input to the network is a volumetric image of (256,256,128). The output is a mask image of the same size.

This U-NET has an encoder architecture with 5 different levels to down-sample the input image. In each level two 3D convolutions are applied with kernel-size (3,3,3). The number of applied filters per convolution changes according to the level: 8, 16, 32, 64, and 128. Each convolution is followed by ReLu activation functions. A batch normalization layer is applied after both convolutions have been performed. A Max-pooling operation is then performed in each level (with exception of the last one). The objective is to downsample the output volumes by half before getting to the next level. Furthermore, the residuals for each level are stored to be concatenated with the corresponding decoder stage levels. The decoder stage is also composed of five levels. An upsample operation is performed in each one, followed by

*The work of L. Cendejas-Zaragoza was supported by CONACyT through a posdoctoral grant (# 647956), and the work of D.E. Rodríguez-Obregón was supported by CONACyT through a doctoral grant (# 787212).

[1,2]L. Cendejas-Zaragoza is with School of Engineering and Science at Tecnologico de Monterrey, and Instituto Nacional de Ciencias Médicas y Nutrición Salvador Zubirán, CDMX, Mexico. email: l.cendejas.zaragoza@tec.mx

[3]D.E. Rodríguez-Obregón, A.R. Mejía-Rodríguez and E.R. Arce-Santana are with Faculty of Sciences, Autonomous University of San Luis Potosi (UASLP), S.L.P., Mexico. emails: diomarrdz@gmail.com, aldo.mejia@uaslp.mx, arce@fciencias.uaslp.mx

[1,4]A. Santos-Diaz is with School of Engineering and Science, and School of Medicine and Health Sciences both at Tecnologico de Monterrey, CDMX, Mexico. corresponding author email: alejandro.santos@tec.mx

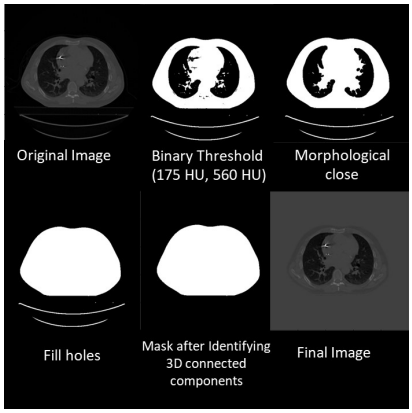

Fig. 2. Body extraction sequence

a transpose 3D convolution layer, ReLU activitation, and a concatenation with the residuals coming from the encoder stages. A normal 3D convolution with a kernel size of (3,3,3) in addition to a ReLU activation is performed afterwards. The convolution operations in the decoder stage use the same number of filters as the corresponding encoder stage at the same level. A final convolutional layer is applied to the upmost level with a sigmoid activation function to create the desired volumetric mask output.

1) U-NET training dataset: An existing lung segmentation dataset was obtained from The Cancer Imaging Archive: Thoracic Volume and Pleural Effusion Segmentations in Diseased Lungs for Benchmarking Chest CT Processing Pipelines (PleThora) [3]. This dataset contained a total of 402 CT scans with their corresponding left and right thoracic cavity segmentations of the lung parenchyma obtained from subjects with diseased lungs.

The provided 402 images, and their associated segmentation masks, were randomly rearranged in training and validation datasets. The training dataset contained 70% (n=281) images, while the validation dataset contained 30% (n=121).

2) Data pre-processing: Lung parenchyma segmentation: Thoracic CT images from the PleThora training and validation datasets were extracted from DICOM files and resampled through a linear interpolation to a (512,512,128) size.

A body extraction algorithm was then performed as follows: First, a binary threshold was applied between 175HU and 750HU intensity values. Then, a 3D morphological close filter with kernel radius (5,5,5) was applied. Next, a binary fill-hole filter was applied. Finally, 3D connected components were identified. Components with a volume of less than two liters were removed; see Fig.2 for a visualization of the body extraction sequence.

3) Data augmentation: PleThora CT images were resampled to a (256, 256, 128) size. Associated segmentation masks were also resampled to the same space. To expand the PleThora training dataset, a data augmentation technique was implemented, which included the following methods: First, a random flip of the x and/or y axes was applied. Then, a 3D random rotation among the x,y, and z axes was performed; the applied rotation varied in a range of (-20, 20) degrees. Next, a 3D random translation in the range of (-5,5)mm among every axis was implemented. Then, the images were normalized in such a way that the intensity values ended with a mean=0 and a standard deviation std=1. Finally, random Gaussian noise was added with a standard deviation in the range of std=(0,0.03).

4) Training the U-NET and k-fold Cross-Validation: The pre-processed PleThora training dataset (n=281 images) was used to train the U-NET by minimizing a cost function based in a Dice Coefficient metric [4]. The training sequence was implemented with a batch size=3, 100 epochs, and an Adam Optimizer with an initial learning rate of 0.005. The validation dataset (n=121 images) was used to evaluate performance of the U-NET training process. An additional k-fold cross validation routine (k=5) was implemented over the PleThora dataset to evaluate overall performance of the network.

The network architecture, training, and cross-validation routines were designed in Keras-Tensorflow [5]

5) U-NET applied to the COVID-19 patient dataset: Once the U-NET was trained with the PleThora dataset, it was used to extract the lung parenchyma for the CT images of the COVID-19 patients. The same pre-processing methodology was applied to these images.

C. COVID-19 lesion volume estimation

Lesion segmentation was achieved using the probabilistic active contours (PACO) algorithm proposed in [6]. The PACO algorithm consists of the minimization of the next functional which guides a contour to the border of the object to be segmented.

$$E[\phi_1, ..., \phi_N, P_1, ..., P_N | I] = \sum_{k=1}^{N} \left\{ \lambda_k \int_L -P_k(V(x)) H(\phi_k(x)) \prod_{\substack{j=1 \\ j \neq k}} (1 - H(\phi_j(x))) dx + \rho_k \int_L |\nabla H(\phi_k(x))| dx \right\} \quad (1)$$

Where V is a volume observed over the voxel lattice L with N disjoint regions. The functions $P_k(.)$ measure how much $V(x)$ belongs to one of the N regions. H is the Heaviside function defined by

$$H(x) = \begin{cases} 1, & \text{if } x \geq 0 \\ 0, & \text{if } x < 0 \end{cases} \quad (2)$$

where $|\cdot|$ is the L_2 norm, and ∇ is the gradient function.

For the first term of this functional, a probability density function (PDf) is used to measure how much voxels belong to each class, and the Heaviside function forces the voxels to be in a single region. The second term is the classical regularization term used in active contours algorithms, which measures the area of the surface that delimitates each volumetric region. The constants λ_k and ρ_k are used to control the influence of each term.

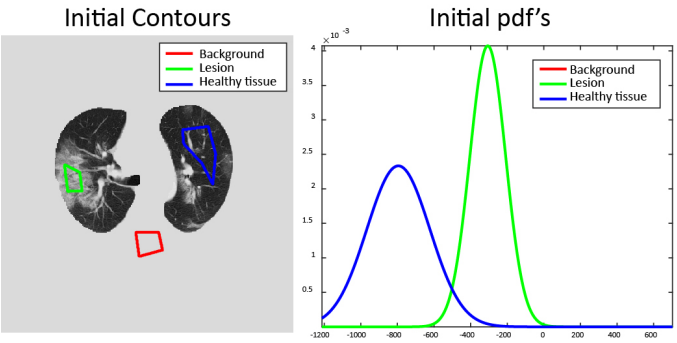

Fig. 3. Initial contours and initial P_k

In order to minimize the functional, PACO uses a two-step minimization process. In the first step PDFs P_k are estimated using the voxels delimited by ϕ_k. In the second step, E is minimized with respect to ϕ_k using the Euler-Lagrange's first variation. These two steps are run repeatedly until achieving the convergence criterion.

For the lesion segmentation presented in this paper, the algorithm was implemented in MATLAB R2020b and the P_k were assumed to be normal distributions. Given the difference in intensity between healthy and diseased tissue, the P_k are not overlapped, leading us to think that estimate the functions with other methodologies, for example Parzen-windows, will lead to similar results.

To initialize the algorithm, three regions of interest were manually delimited on one of the volume slices. This slice was selected from the middle slices, searching for one with a large lesion area. An example for one patient, of the initial contours is shown in Fig. 3 along with the initial probability distribution functions, the green contour corresponds to the lesion, blue contour to the healthy tissue, and red to the background. Notice that the background class have a PDF that approximates a Dirac delta and is not shown in the figure.

For this application, the λ_k parameters were set to 1 and the ρ_k to 0.5. The iterative process was performed for 100 iterations. The Heaviside, and Dirac delta functions were replaced as in [6] for numerical implementation.

III. RESULTS

A. Lung parenchyma segmentation

Fig. 4A shows the result of applying the lung parenchyma segmentation workflow based in the U-NET in a CT image of the PleThora validation dataset. The first image shows the ground truth. While the second shows the predicted mask after using the U-NET architecture.

The k-fold cross-validation (k=5) results for the U-NET are presented in Table I. These results were obtained when using the PleThora dataset. The average Dice Coefficient in the training datasets was 0.9392, while it was 0.9423 in the validation datasets. The previous numbers indicate that the network performance is optimal for this segmentation task: A Dice Coefficient equal to 1.0 will indicate a perfect segmentation according to the provided masks [4].

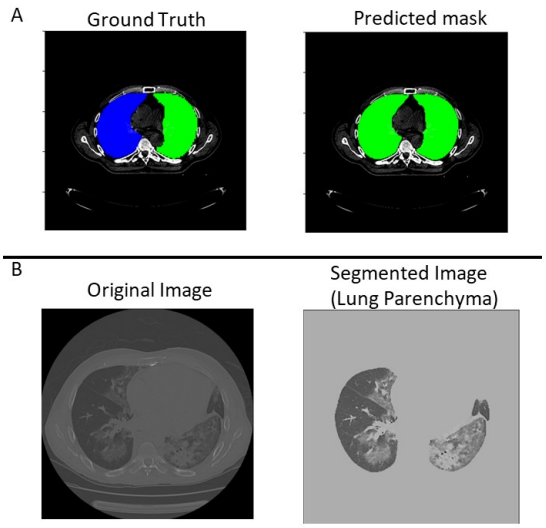

Fig. 4. Lung parenchyma segmentation (U-NET performance)

TABLE I
U-NET K-FOLD CROSS VALIDATION

Fold	Training Loss	Training Dice Coefficient	Validation Loss	Validation Dice Coefficient
1	0.0649	0.9351	0.0688	0.9312
2	0.0680	0.9320	0.0583	0.9417
3	0.0499	0.9501	0.0442	0.9557
4	0.0634	0.9366	0.0608	0.9392
5	0.0578	0.9422	0.0565	0.9434
Average	0.0608	0.9392	0.0577	0.9423

Fig. 4B shows the results of applying the lung parenchyma segmentation workflow to a CT image for a patient with COVID-19. As expected, the U-NET network delivered acceptable qualitative results. All CT images for the 50 COVID-19 patients were processed using the U-NET. All lung segmentations were visually inspected.

B. COVID19 Lesion Volume Estimation

Fig. 5 shows the result segmentation using PACO over the parenchyma image in four different slices from different patients. The slices appear in different rows in the figure, the first image of each slice is the parenchyma image which is the input image for PACO. The second image is the multi-class output segmentation, with the white color representing the lesion mask segmented. The third one is the lesion mask overlapped with the parenchyma; this image allows to observe qualitatively, the regions where the algorithm makes a correct segmentation or where it fails.

The algorithm segments correctly the diseased area. Nevertheless, it also includes other regions in the lesion mask, such as vasculature, airways, and some regions that are not parenchyma but were delimited by the U-NET.

All 50 patients were categorized in two groups according to their clinical outcome (survived or died). Our pipeline was applied to their CT images. The percentage of lesioned

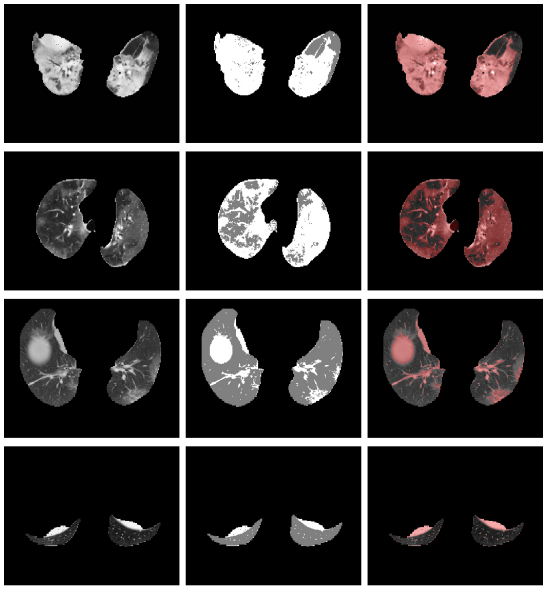

Fig. 5. COVID19 lesion segmentation with PACO.

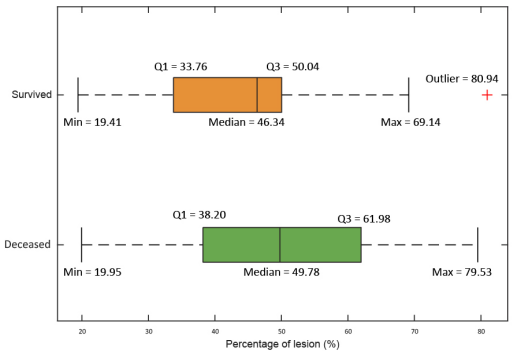

Fig. 6. Boxplots for percentages of lesion in survived and deceased patients.

tissue against healthy tissue was obtained per individual. This value was calculated by counting and comparing the number of voxels in the output masks. Fig.6 shows a comparison of the lesion-related percentages observed per group (survived or died).

IV. DISCUSSION

We proposed a workflow consisting of two steps to obtain a COVID-19 lung lesion segmentation from CT images. The first step consisted in isolating the lung parenchyma using a convolutional neural network following a U-NET architecture. The second step consisted in applying the PACO algorithm to create the final lesion mask. Visual inspection of the output images shows that the proposed pipeline successfully identifies lesion areas. However, as it can be seen in Fig.5, other areas which are non-related to the lesion have been marked by the output mask. Specifically, areas within lung airways and lung vasculature. Other mislabeled regions include tissue that is not part of the lung parenchyma. These regions are selected since the U-NET output has errors in these areas. That is, the U-NET classifies erroneously several voxels of the original CT as being part of the lung-parenchyma. This error is propagated down our pipeline which affects the performance of the final lesion segmentation performed by PACO, which is a contour-based algorithm that is limited by an initial selection of classes. We used three classes for the PACO algorithm.

The k-fold cross validation results suggest that the chosen U-NET architecture is optimal for the lung-parenchyma segmentation task. However, these results were obtained in the PleThora dataset which include CT images whose imaging sequences differ from the ones obtained for our COVID-19 patient cohort.

The performance of the U-NET in COVID-19 CT images can be increased by visually inspecting every lung parenchyma segmentation, and manually correcting the masks. Thus, creating a COVID-19 training dataset for segmentation, which can be used for benchmarking future automatic segmentation pipelines. The U-NET can be re-trained using this new dataset.

The median of the percentage of the lesion in the deceased group appears on the third quartile of the survived group, making this metric an important feature to evaluate COVID-19 patient prognosis.

V. CONCLUSIONS

Our two-step pipeline successfully segmented COVID-19 related lesions in CT images, with the exception of some mislabeled regions including lung airways, vasculature, and tissue not related to lung. Our workflow can be applied to derive a volume metric related to lesion to evaluate COVID-19 patient prognosis.

ACKNOWLEDGMENT

Authors acknowledge the Critical Care and Radiology Departments of the Instituto Nacional de Ciencias Médicas y Nutrición Salvador Zubirán (INCMNSZ), for their support on the present research.

REFERENCES

[1] WHO, "World health organization coronavirus (covid-19) dashboard." [Online]. Available: https://covid19.who.int/

[2] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2015, vol. 9351, pp. 234–241.

[3] K. J. Kiser, S. Ahmed, S. Stieb, A. S. R. Mohamed, H. Elhalawani, P. Y. S. Park, N. S. Doyle, B. J. Wang, A. Barman, Z. Li, W. J. Zheng, C. D. Fuller, and L. Giancardo, "PleThora: Pleural effusion and thoracic cavity segmentations in diseased lungs for benchmarking chest CT processing pipelines," *Medical Physics*, vol. 47, no. 11, pp. 5941–5952, nov 2020.

[4] S. Jadon, "A survey of loss functions for semantic segmentation," in *2020 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB)*. IEEE, oct 2020, pp. 1–7.

[5] F. Chollet *et al.*, "Keras," https://keras.io, 2015.

[6] E. R. Arce-Santana, A. R. Mejia-Rodriguez, E. Martinez-Peña, A. Alba, M. Mendez, E. Scalco, A. Mastropietro, and G. Rizzo, "A new probabilistic active contour region-based method for multiclass medical image segmentation," *Medical & biological engineering & computing*, vol. 57, no. 3, pp. 565–576, 2019.